Use of directly obtained GPS velocity in the computation of winds and currents instead of a velocity derived from GPS position

Jorge F Garcia, LT USN

Abstract

To investigate the error in oceanographic and meteorological measurements that require a computation of ship's velocity and to explore the possible reduction in said error by using GPS velocity data instead of GPS position, a comparison of three time series was performed. The first time series contains GPS velocity as sampled by the ASHTECH receiver at a rate of 1 Hz. The second time series contains ships velocity as used by the Acoustic Doppler Current Profiler (ADCP) with a 'sample' rate of 1/300 Hz. The third time series contains ship's speed (not velocity) as recorded by Serial ASCII Interface Loop (SAIL) and then used compute true wind with a 'sample' rate of 1/48 Hz. Due to the different sampling rates and the absence of a standard time reference frame, most of the work involved manipulating the data as required for an exact correlation of time and an overlap of good data from all three time-series. Due to the high sample rate and accuracy of GPS data, it had to be processed to remove the high-frequency pitch and roll components that - due to averaging - are not present in ADCP or SAIL data. Then the data was compared and conclusions drawn.

Introduction

Modern oceanographic vessels rely upon the Global Positioning System in order to accurately measure winds and currents with shipboard instrumentation. The ship's velocity obtained from GPS needs to be subtracted from the measurements in order to transform relative measurements into earth-reference ones. Currently the wind and current calculations used on data gathered by the **RV Point Sur** are based upon a ship's speed that is calculated by averaging GPS positions over a period of time.

GPS can provide a user anywhere in the world with a position that has a six-meter spherical error probability. If the user is within operating range of a Differential GPS station, the error can be reduced down to a 2 centimeters SEP. This position-function of GPS has been widely exploited for scientific and military purposes. If the ship's velocity was constant, then the random error could be essentially removed by averaging over a long enough period of time. However, due to the variable effects of numerous internal and external factors, ship's velocity cannot be treated as a constant. Therefore, some information - and accuracy - is necessarily lost in the averaging process.

Ignored by most, however, is the fact that GPS can provide several other products. Some GPS receivers - and in particular the ASHTECH receiver on board **RV Point Sur** - can be configured to extract heading, pitch, roll and velocity *directly* from GPS Doppler information. These are separate products and are not derived from GPS position.

Therefore, one at least apparent advantage of using GPS velocity instead of position is that information is not lost in the averaging process. Another, perhaps more important advantage is that GPS velocity accuracy is better than 10 centimeters per second, *without* operating in differential mode. It is as good as 1 centimeter per second with the aid of Differential.

With this in mind I decided to attempt to quantify the increased accuracy that might be obtained from using the velocity product directly from GPS, and to produce a qualitative assessment of the desirability to switch to this source of data.

Difficulties

Due to its volume, the primary source of difficulty was the manipulation of the GPS data. 78500 seconds (slightly less than a day) were stored as a thirty-megabytes file in a **.txt** format (FIG.1).

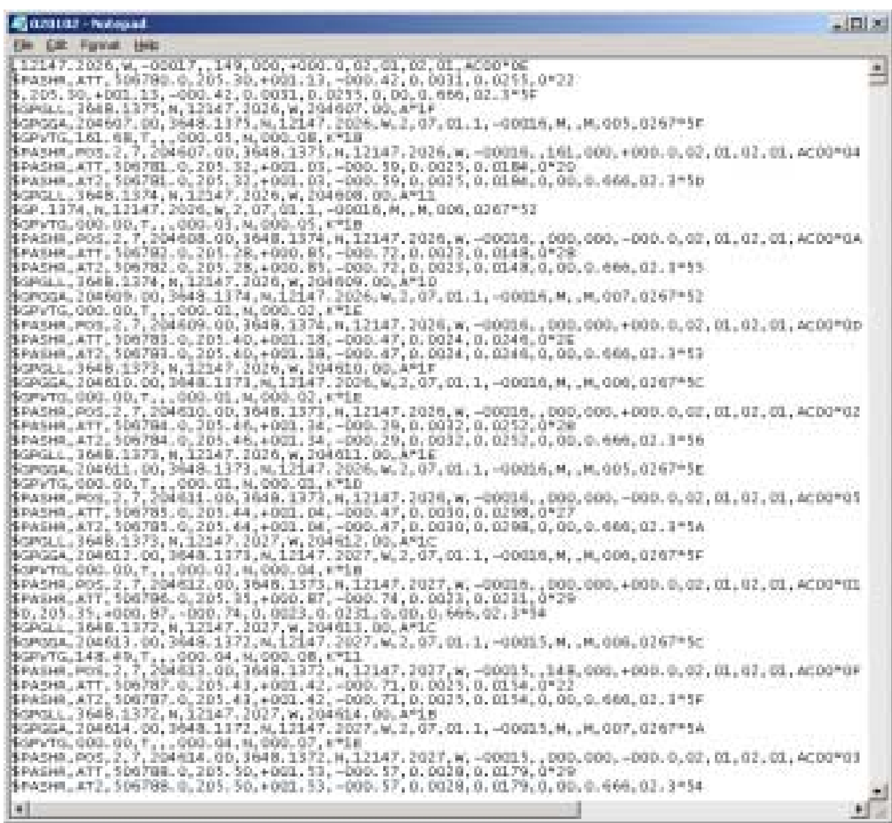

FIG. 1. This is an example of a portion of the output from ASHTECH in .txt format. Note the different fields containing different kinds of data, some of which is redundant.

Part of the problem presented by the volume of the data is that the large data-stream apparently exceeded the efficient storing rate capacity of the computer. This introduced a lot of noise bits into the file. NMEA2ESL 3.4, a program designed by Prof. James R. Clynch to transform the .txt file into ESL (a Matlab-readable format) discovered 160235 line errors in 471938 lines (REF 1).

These errors are most significant in the context of this work in that some of them cause the time to be misread out of the file. Thus the first challenge - after the jumps in time were explained - was to correlate the GPS time series to the same reference frame as the others (FIG.2).

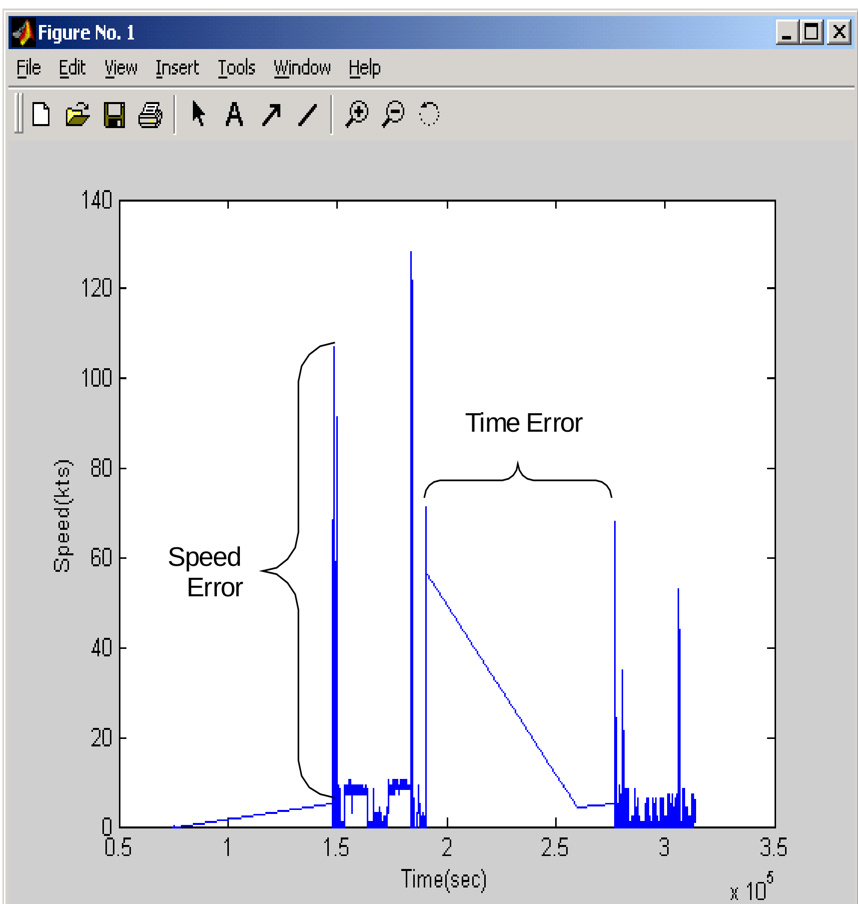

FIG.2. The entire time-series of speed calculated from GPS velocity showing some effects of the noise. Although these resulted in very large errors, other effects may not be as obvious.

Instead of removing the time-noise and splicing the clean sections together, I opted for simply using a section of the data that had no *obvious* noise. The section selected spanned slightly over 25000 seconds (approximately eight hours) of data.

As soon as the noise causing the spikes in the speed was thus removed, a second problem became evident. The ship appeared to have been doing only exact increments of speed, in knots (FIG.3).

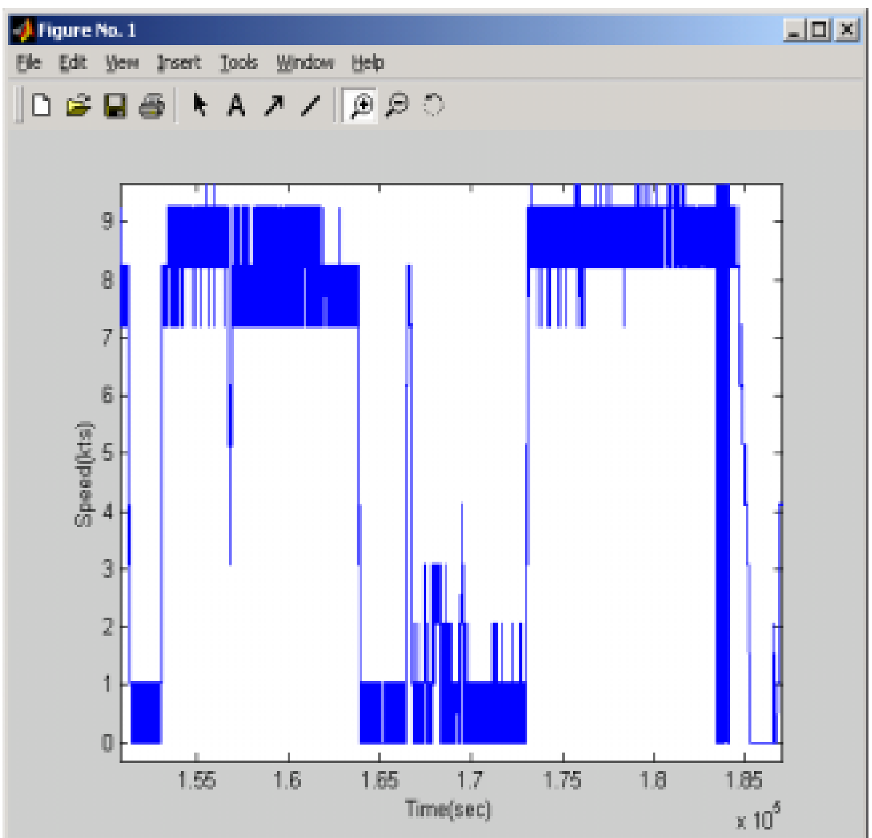

FIG.3. There is obviously a problem here. The data is supposed to have a resolution on the order of centimeters per second, but clearly - somewhere along the process - it has been rounded off to the nearest knot. The problem turned out to be that the receiver was configured in *aircraft mode.*

It developed that the ASHTECH receiver can be configured to output data in a variety of formats. Some of these formats are better suited for aircraft than for shipboard use. This case proved to be an example of such configuration. The receiver had been set to output velocity data to the nearest knot in order to allow for a useful display of a velocity that generally ranges in the hundreds of knots.

Fortunately, the ASHTECH will output some redundant information in different *fields*. While the information in the *POS* field forsook available GPS accuracy for practical reasons (assuming the **RV Point Sur** was an aircraft), velocity information was also available in the *VTG* field. In *this* field, the receiver had been configured to output velocity data in an **X.XX** (m/s) format.

Prof. Clynch edited the data in the .txt file such that the *POS* field became unusable. The NMEA2ESL program then ignored these and went to the next available field in search of the required information. Once the good velocity data was available, E and N velocities (m/s) were converted to speed (kts) in order to have the data in a more intuitively accessible format.

Procedure

Now that the data was good, the next step became apparent: the necessity for extracting high frequency component which rides over ship's speed and which is due to pitch and roll (FIG.4). The magnitude of this oscillation was slightly less than two knots.

Using direct samples from this data would potentially introduce a huge error in ship's speed.

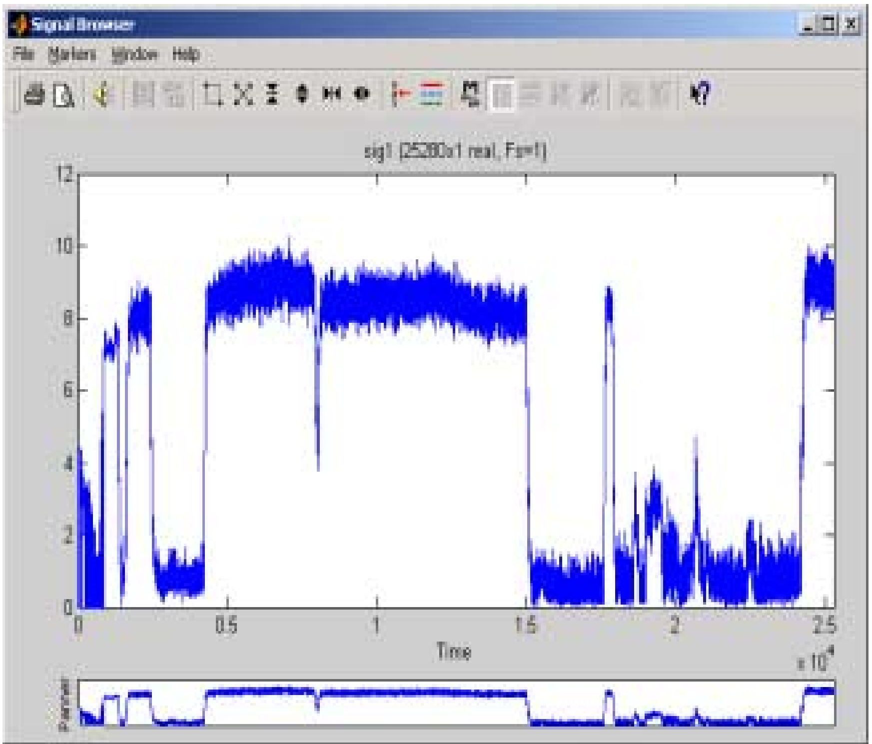

FIG. 4. This is the time-series of ships speed obtained from clean and good GPS velocity data (Approximately 8 hrs). Note the effects of pitch and roll.

A close look at the speed oscillation revealed a fairly steady frequency, a period of approximately 3.5 seconds. This seems t correspond closely with a previously measured roll period for the Point Sur of approximately 8 seconds (Due to the fact that all speeds are positive, a roll period is twice as long as the period of the change in speed). Note also that since this information is being extracted from speed, without referencing ship's heading it can't be determined whether the oscillation observed is due to roll, pitch or yaw. Finally, note that if we assume this oscillation to be due to roll alone, then only the *minimums* in the oscillation would correspond with ship's speed (FIG.5).

Three other products directly available from GPS are the heading, pitch and roll. A comparison of these factors with the characteristics of the measured speed oscillations was therefore possible (FIG.6 and FIG.7). Consideration was given to the possibility of using these products in order to extract the effects of pitch and roll from the velocity. This would require that the GPS data be left as East and North velocities instead of speed. The pitch and roll, measured in degrees, would be converted to velocity by using the height of GPS antenna above the ship's metacenter. Then the ship's heading would be used to shift the velocities from east-west to atwartship-foren'aft coordinates.

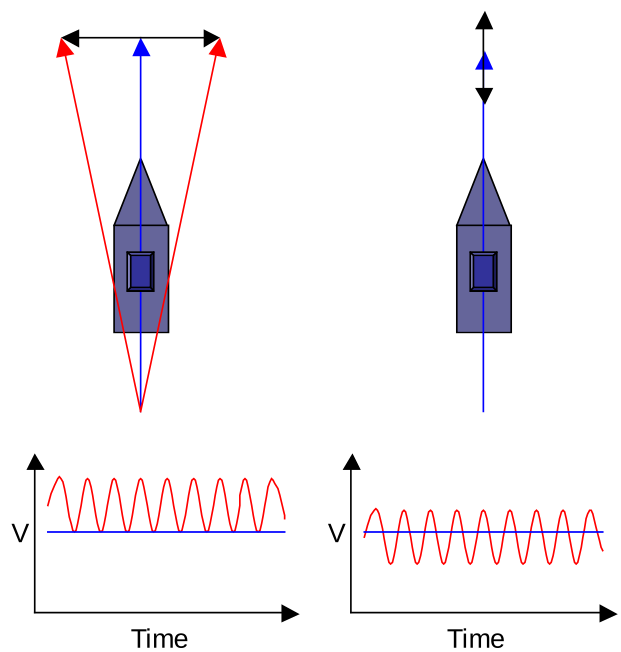

FIG.5. Due to the fact that speed is a scalar quantity, a roll (perpendicular to it) can only increase it. Note also that the ship's speed is *true* twice on every roll cycle (as the ship reaches the extreme of its roll and starts to turn around). Thus the Oscillations in speed have twice the frequency of the roll.

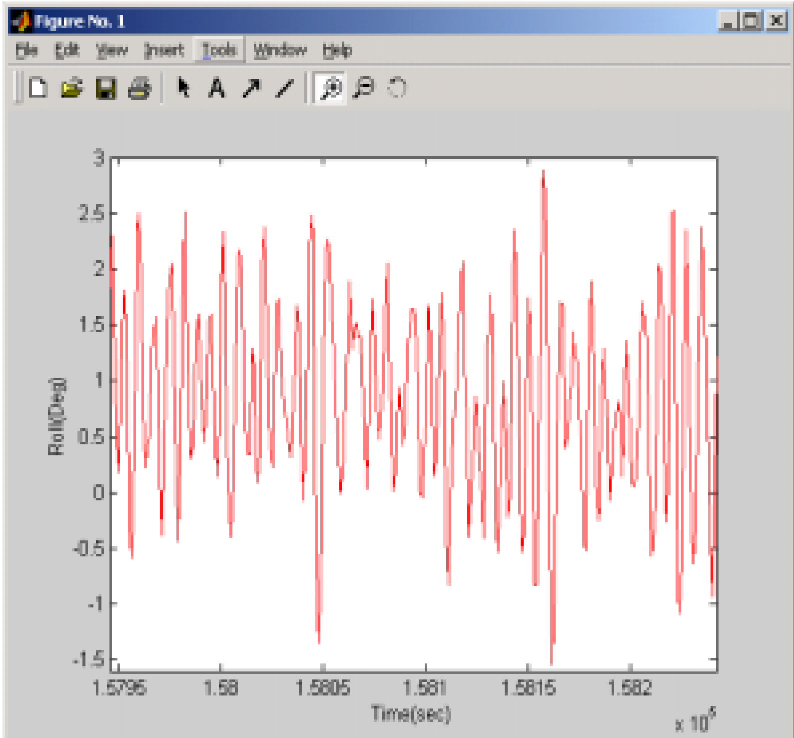

FIG.6 Ship's roll obtained directly from GPS. Note the regularity of the frequency. Period is approximately seven seconds, a close match to the oscillation in ships speed. Note that there is about one degree bias. The relatively smooth changes in magnitude make the roll look like an *amplitude-modulated signal*

The problem with attempting to use pitch and roll from GPS to remove their effect from velocity is that the noise has affected the pitch and roll time-series in the same way that it affected the velocity time-series (by causing time skips), but at different points (since the data comes from different fields). Due to the relatively high frequency of these oscillations, a shift of one second (one missed data point) would render the output of a transformation in coordinates useless. Since it is impossible, or at least very difficult to verify that both time-series are coherent for *every* data point out of approximately 25000, I decided to attempt something else.

.

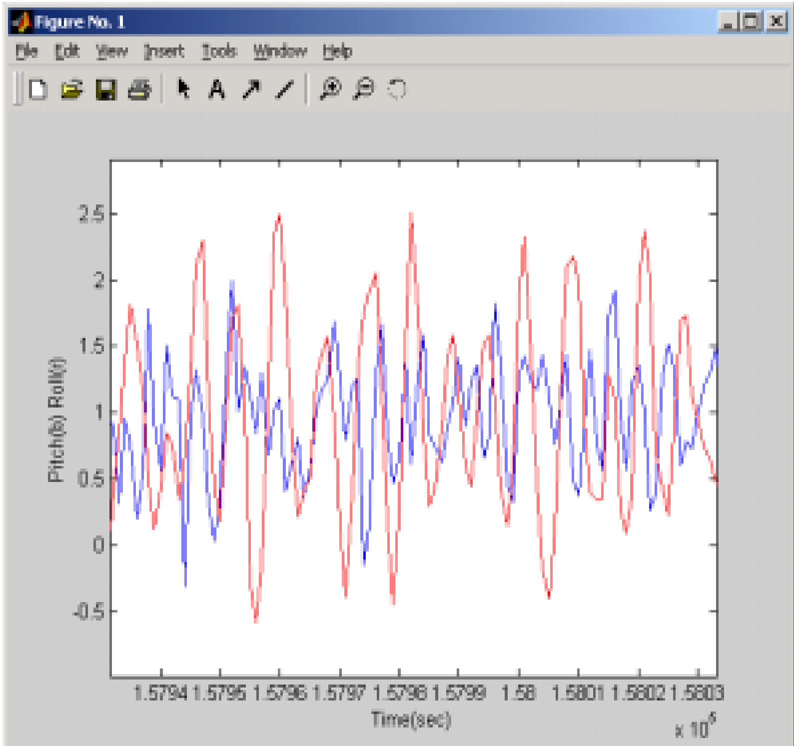

FIG.7 Pitch from GPS is compared with roll. Note the higher frequency, reduced regularity, and reduced magnitude of pitch. It makes sense since the ship is more stable foren'aft than athwartships, and the pitch frequency is more readily influenced by ship's course (relative to the seas)

An attempt to unscramble the pitch and roll using the method described above would require a lot of data manipulation and would potentially yield nothing. Instead of this approach, a variety of filters were applied to the data (FIG. 8). It was possible that in using this method, the resulting filtered data would have a positive bias (due to the previously discussed effect of roll wherein port and starboard roll *both add* to ship's speed). However, a simple visual comparison between the filtered data and the data obtained from ship's speed would give an indication of the magnitude of this bias. It could then be corrected for.

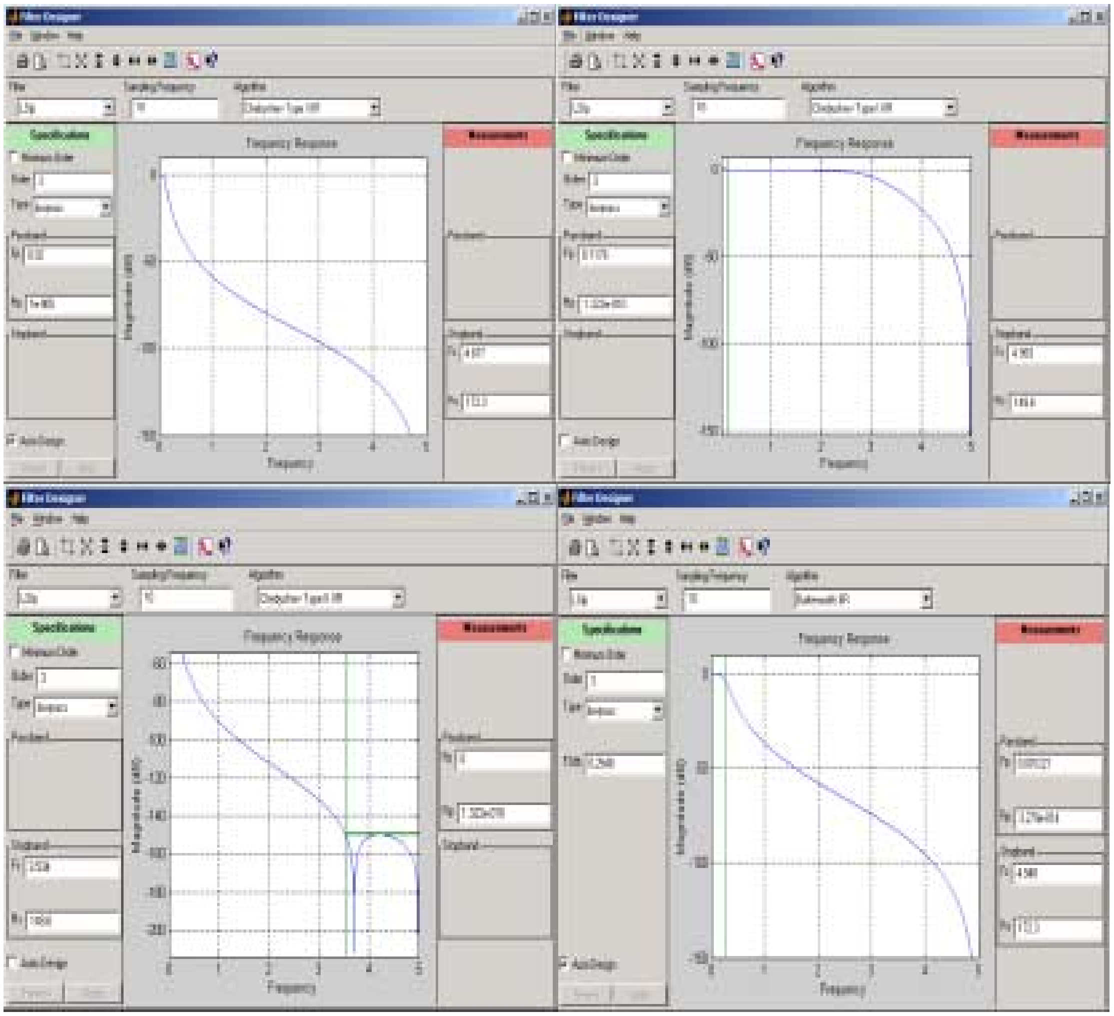

FIG. 8 The frequency response curves of the various filters tested. The filters where designed using Matlab SPTOOL (REF 2), an interactive signal processing tool. Note that all of them are designed to attempt to block signals of frequencies of about 3.6 Hz and greater.

The resulting filtered data was then closely inspected and compared to verify that in fact the high frequency oscillation was being removed without "damaging" the low frequency ship's speed. The best performing filter was thus selected by visual inspection of the plots (FIG. 9).

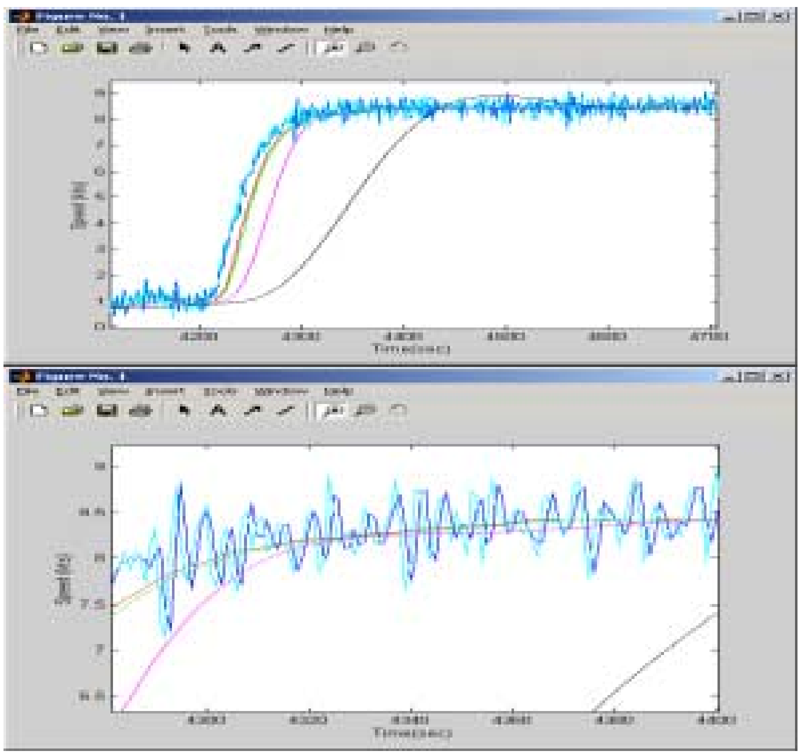

FIG. 9 A visual comparison of the performance of five different filters. The actual data is in dark blue. The filter depicted in light blue does not succesfully remove the high frequency oscillations while the black and magenta remove too much. The red and green provide the best fit.

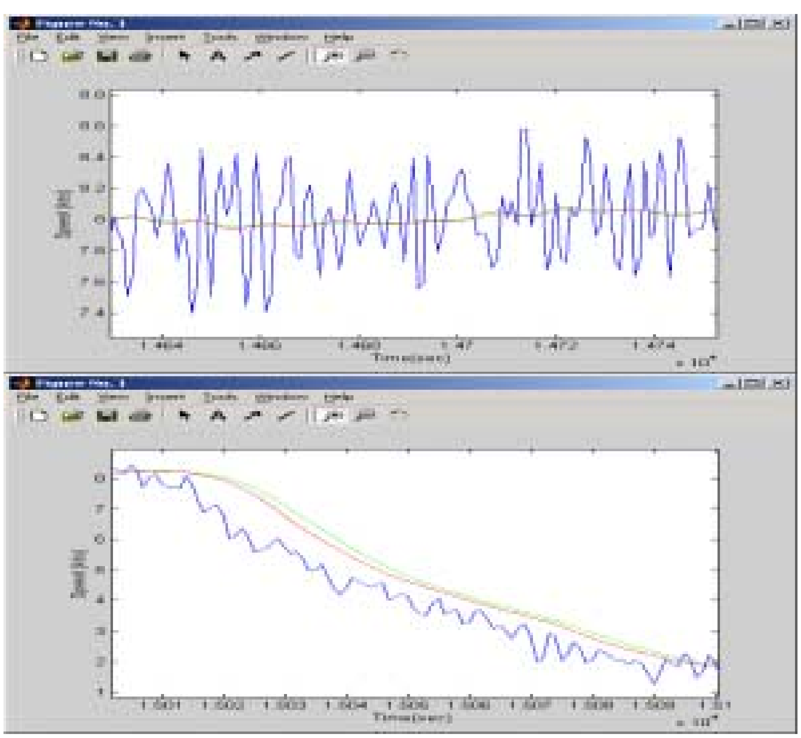

FIG.10 The two best performing filters. Note that if the actual ship's speed corresponds to the point in between peaks, then the error from the filter is less than 0.1 kts when the ship is at constant speed. However the error grows to slightly over one knot during acceleration periods.

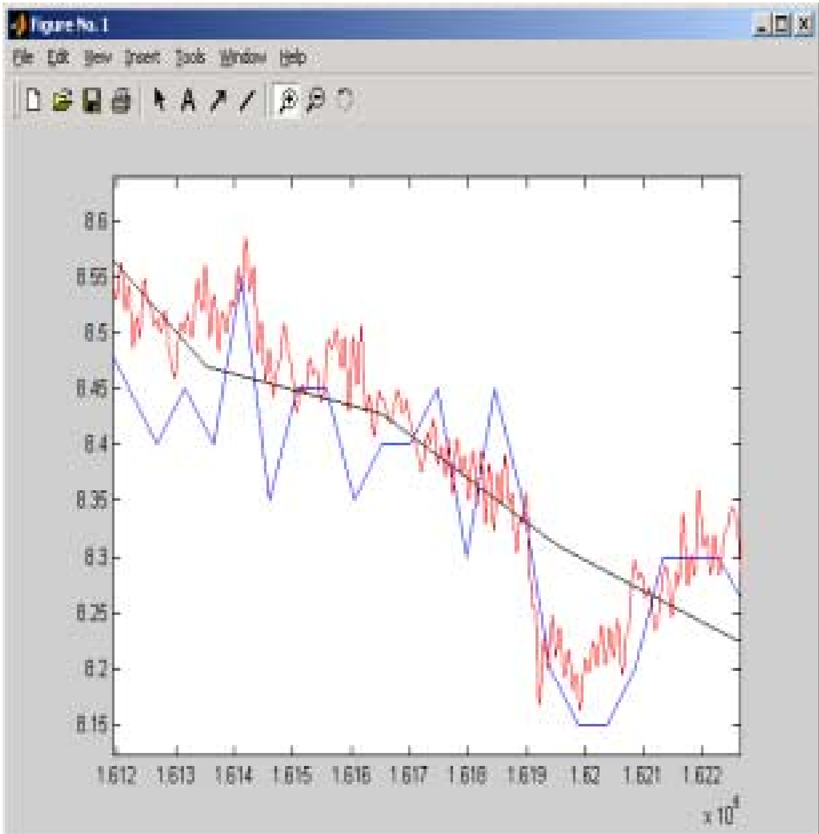

FIG.11 The output from the Butterworth IIR filter shown in the lower right hand of FIG. 9 was selected. Shown here, the filtered speed data (red) now almost ready to be compared with ADCP (black) and SAIL (blue). Note that SAIL data is much closer to filtered GPS speed than ADCP. Note also that no bias appears to have emerged from the filtering process.

The final step consisted of decimating the filtered GPS data in order to match the much slower sample rates of ADCP and SAIL. Once the time series were of the same size, the difference was taken and plotted after having been squared and square rooted to ensure that the resulting statistical data referred to the *magnitude* of the difference between data points.

It was clear that the performance of the filter was severely degraded during acceleration periods (FIG 10). Therefore two sets of statistical analyses were performed, one for steady ship's speed, one for the whole time-series.

	SAIL Overall	SAIL Constant Speed	ADCP Overall	ADCP Constant Speed
Min	0.0001	0.0003	0.0002	0.0022
Max	3.597	0.3544	1.766	0.7399
Mean	0.09207	0.04728	0.2334	0.08514
Median	0.05512	0.03691	0.08126	0.04593
Std	0.1808	0.04745	0.3947	0.1409
Range	3.597	0.3541	1.765	0.7377

TABLE.1 Statistics quantifying the magnitude of the difference between the filtered GPS speed and the speed used by ADCP and SAIL. Considering the filtered GPS speed to be the most accurate, the mean and standard deviation are an indication of accuracy relative to this standard. Note that SAIL data has a much closer fit than ADCP. It is also important to consider the magnitude of the quantities being ultimately measured. Since winds are approximately one order of magnitude greater than currents, the error incurred by ADCP is actually more significant than the statistics imply.

Conclusions

Shortly before concluding the data analysis section of this project, I found out that SAIL does not actually calculate velocity by using GPS position over time. Instead, SAIL does in fact use the very product whose better accuracy - and consequently, its desirability - this paper was meant to address. Instead of being a setback, this unexpected development serves to provide a fairly robust authentication of my work.

SAIL actually uses an average of GPS velocity over thirty seconds. This time period is enough to eliminate the effects of pitch and roll, however it is short enough to retain almost all the information on ship's speed. Therefore SAIL data is actually very accurate. The fact that it closely correlates with the filtered speed data that I extracted from GPS velocity confirms my assumptions (FIG. 11).

On the other hand, ADCP does use position (sampled at approximately 1/6 Hz) over a five minute period to perform its calculation of velocity. The issue is addressed by the literature (REF 3):

> "[If bottom tracking is not available] GPS is the best choice for general purpose accuracy. Some GPS systems can directly use GPS Doppler information to obtain ship's velocity. The use of this information can reduce short-term error in ship's velocity by about an order of magnitude. GPS data that use the GPS Doppler information are normally much more accurate than the ADCP data, so the GPS will not affect the current profile accuracy."

Which indicates that the designers are aware of the advantages presented by using GPS velocity, but realize that given the limitations of ADCP, no accuracy is gained. The passage appears to indicate, however, that GPS velocity *ought* to be used. This would fit with the claim in the same reference that ADCP bias (the error which cannot be removed by averaging) is on the order 0.5 to 1 cm/sec.